REMARKS

This paper is responsive to the Office Action dated November 19, 2004. Claims 21-55 were examined.

Rejections Under 35 U.S.C. §102

The Examiner rejected claims 21-55 under 35 U.S.C. §102(b) as being anticipated by U.S. Patent No. 5,168,564 issued to Barlow (hereinafter "Barlow"). Applicant respectfully traverses all rejections.

Applicant respectfully submits that the Office has misunderstood Barlow's cancellation mechanism. The Office presumes that "the lock indicator, or mechanism can be canceled after it is set once it is determined that the command using the resource that is locked is invalid, therefore the resource is being locked before the command has been determined to have hazards." Barlow never discloses or suggests speculative locking of a resource, determining whether hazards exists, and especially does not disclose speculatively locking a resource prior to determine whether a hazard exists.

Barlow's mechanism couples a cancel command with either a read set lock command or a write lock reset command to respectively override either the read operation or the write operation. When the cancel command accompanies a read set lock command, the read set lock command is executed as a lock command (col. 8, lines 38 – 58). When the cancel command accompanies a write lock reset command, then the write lock reset command is executed as a lock reset command (col. 8, line 59 – col. 9, line 52).

Furthermore, the cancel command does not occur prior to determination of whether a hazard exists or determination of whether an operation is invalid. Indeed, the cancel command I issued with a read set lock command after "when it is discovered that it becomes desirable only to modify the state of the memory lock bit without disturbing the rest of the memory subsystem" (col. 8, lines 33 – 35). "When the CSS subsystem wants to reset the state of the memory lock bit, it generates a write lock reset command accompanied by a cancel command" (col. 9, lines 59 – 62). It is clear that Barlow neither discloses nor suggests "speculatively locking a resource...prior to determining whether a hazard exists" as recited in claim 21, and similarly

recited in claims 32 and 45, and does not disclose or suggest "wherein respective processing cores are adapted to lock the resource in response to respective accesses by respective first instructions prior to determining whether a hazard exists between the respective accesses and the second instruction" as recited in claim 33.

The Office relies upon the same arguments as proffered for claims 21, 32, and 45, to reject claim 45. However, as already stated, Barlow discloses issuing the cancel command ***after*** determining a desire to alter state of the memory subsystem. Claim 42 recites "[a] processor to speculatively dispatch a load operation to a cache unit prior to determining whether read-after-write hazards associated with the load operation are present." Barlow does not disclose or suggest a processor to dispatch a load operation to a cache unit, and especially does not disclose such a processor to dispatch prior to determining whether a read-after-write hazard is present. Applicant respectfully requests guidance to understand how Barlow's cancel mechanism discloses or suggests a processor as recited in claim 45.

With regard to claims 22, 34, and 46, the Office relies upon the following 2 sections:

> The cancel command of the present invention finds use in a variety of different types of operations. It becomes useful in issuing diagnostic commands identifiable via bus signal BSYELO+ which is described in the Appendix. Such diagnostic and cancel commands can be issued to a memory subsystem or I/O subsystem where it is necessary to set or reset reconfiguration state registers or other elements useful in diagnosing system faults. In this situation, each such command with cancel, upon being decoded, causes the required change of state without causing the memory or I/O subsystem to initiate any cycle of operation. (col. 7, line 60 - col. 8, line 3)
>
> The lock address register 100-210 monitors the addresses being forwarded to the system bus 11 and saves the address associated with last lock operation performed by the CSS unit. The physical address register 100-214 is involved in the overall recovery process. If a failure occurs in the pipeline stages for any reason, the address contents of this register indicate to the software the last physical address or I/O channel number that was being used by the CSS unit in connection with accessing system bus 11. (col. 5, lines 9 - 18)

The Office also refers to Figure 4A of Barlow, which depicts timing diagrams with bus cycles. Neither the figure nor these sections disclose or suggest a trap stage, and especially do not

disclose or suggest "wherein the locking is performed prior to the first instruction entering a trap stage of an instruction pipeline" as recited in claim 22, and similarly in claims 34 and 46. The Office basis the rejection of these claims on the sole argument that a fault is a trap, disregarding the actual claims.

With regard to claims 23, 35, and 47, the Office asserts that Barlow teaches an atomic instruction as recited in the claims. The Office refers to the disclosure of a read-modify-write operation in the background section of Barlow. Nothing in this section discloses an atomic operation including a portion to lock a resource and a portion to unlock the resource. In fact, Barlow actually discloses 2 separate instructions to implement the RMW operation: the read set lock command and the write reset lock command.

With regard to claims 24, 36, and 48, the Office again presumes the disclosure of subject matter that is not found in Barlow. The Office states "Barlow has taught wherein the hazard includes a read-after write hazard" and cites the following section:

> Most access difficulties are encountered when a processing unit is required to perform an instruction sequence specifying a read-modify-write (RMW) operation. In an RMW operation, one processing unit fetches data from a location in memory, performs an operation on the data contents of the location and writes the modified data back into the original memory location.
>
> One way to prevent more than one processing unit from performing a RMW operation on the same memory location, an interlock read instruction is utilized. This involves the use of a lock indicator device, which is set during the read portion of an RMW operation to prevent access to a specific memory location, and is reset after the write portion of the RMW operation is completed. If a second processing unit should attempt to access the same memory location to perform an RMW operation, the memory subsystem will send a busy signal indicating that the memory location is in use.

There is no disclosure of a read-after write hazard in the section relied upon by the Office or in any other section of Barlow.

With regard to claims 25, 37, and 49, the Office rejects these claims based on Barlow's disclosure of address generation at col. 4, lines 35 – 51 and disclosure of saving an address associated with a last lock operation performed by a CSS at col. 5, lines 9 – 19. However,

Barlow fails to disclose or suggest "locking the resource during an effective address calculation stage of an instruction pipeline" as recited in claim 25, and similarly in claims 37 and 49.

With regard to claims 28, 40,43 – 44 and 52, the Office refers to Barlow's disclosure of a RMW operation. The Office states that "the resource is locked at the read portion and reset after the write portion of the operation – after the write portion of the operation, the process is complete and therefore leave the pipeline." However, there is not indication of any disclosure or suggestion in Barlow for the actual limitations of the rejected claims.

With regard to claims 30 – 31, 41, and 54 – 55, Applicant respectfully submits that the Office has erroneously interpreted Barlow as already discussed above.

Applicant respectfully submits that none of the claims are disclosed or suggested by Barlow, or any other art of record. The independent claims and the dependent claims are allowable at least because of the reasons given above. The dependent claims are also allowable at least because they depend from corresponding ones of the allowable independent claims.

Conclusion

In summary, claims 21 – 55 are in the case. All claims are believed to be allowable over the art of record, and a Notice of Allowance to that effect is respectfully solicited. Nonetheless, if any issues remain that could be more efficiently handled by telephone, the Examiner is requested to call the undersigned at the number listed below.

CERTIFICATE OF MAILING OR TRANSMISSION

I hereby certify that, on the date shown below, this correspondence is being

☐ deposited with the US Postal Service with sufficient postage as first class mail, in an envelope addressed to Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

☐ facsimile transmitted to the US Patent and Trademark Office.

Steven R. Gilliam _____ Date _____

EXPRESS MAIL LABEL: EV 671081794 US

Respectfully submitted,

Steven R. Gilliam, Reg. No. 51,734
Attorney for Applicant(s)
(512) 338-6320
(512) 338-6301 (fax)